

Information Science & Technology Institute – Rapid Response: D-Wave Effort Debrief

Agenda

Los Alamos National Laboratory
Jemez and Cochiti Rooms
2nd Floor of the Research Library

Thursday, October 6, 2016
12:45pm – 5:30pm

- 12:45 **Welcome, Logistics**.....Stephan Eidenbenz (ISTI)
- 1:00 **Efficient Combinatorial Optimization Using Quantum Computing**..... Hristo Djidjev
- 1:20 **Quantum Connect-the-Dots** Zach Baker
- 1:40 **Graph Partitioning using the D-Wave for Electronic Structure Problems**..... Sue Mniszewski
- 2:00 **Generative Modeling for Machine Learning on the D-Wave**..... Sunil Thulasidasan
- 2:20 **Learning Sparse Representations for Object Classification Using the Quantum D-Wave 2X Machine**..... Garrett Kenyon
- 2:40 **Phylogenetics through Quantum Annealing**..... Marcus Daniels
- 3:00 **BREAK**
- 3:20 **Ising Simulations on the D-Wave QPU** Mike Rogers
- 3:40 **D-wave Quantum Computer as an Efficient Classical Sampler** Andrey Likhov
- 4:00 **Challenges and Successes of Solving Binary Quadratic Programming Benchmarks on the DW2X QPU**..... Coffrin Carleton
- 4:20 **Topological Sphere Packing on the D-Wave** David Nicholaeff
- 4:40 **Quantum Uncertainty Quantification for Physical Models using ToQ.jl**..... Dan O'Malley
- 5:00 **DISCUSSION** Lead: John Sarrao (ADTSC), Mark Anderson (ASC)
- 5:30 **ADJOURN**

Efficient Combinatorial Optimization Using Quantum Computing Hristo Djidjev

Abstract:

We develop algorithms for solving NP-hard optimization problems suitable for QUBO/D-Wave implementation. The two NP-hard problems we focus on are the maximum clique problem, whose objective is to find a clique (a subgraph with edges between all pairs of vertices) of maximum size, and two variations of the graph partitioning problem, whose objective is to divide a graph into two parts of equal number of vertices by cutting a minimum number of edges. We study scalability/accuracy issues and ways to mitigate them and compare the performance of the quantum annealing implementations with several algorithms for classical computers.

Quantum Connect-the-Dots Zach Baker

Abstract:

Constrained Shortest Path Estimation on the D-Wave 2X:

Accelerating Ionospheric Parameter Estimation Through Quantum Annealing.

Shortest path computations are a general purpose solution to many problems, but high degrees of connectivity do not map well to the limited connectivity of the D-Wave 2X. By re-casting the problem as a series of 1-of-n choices linked by weighted connections, the shortest path problem is mapped to the quantum machine in a practical and useful way. This approach is demonstrated with an ionospheric parameter estimation problem.

Graph Partitioning using the D-Wave for Electronic Structure Problems.....

S.M. Mniszewski, C. F. A. Negre, H. Ushijima-Mwesigwa

Abstract:

Graph-based methods are currently being applied to electronic structure problems for quantum molecular dynamics (QMD) simulations. Generating the density matrix as part of a timestep from many small sub-matrices (or sub-graphs) has been shown to be equivalent to more traditional methods (such as diagonalization). We have explored relevant graph partitioning/clustering methods and implementations that run on the D-Wave, 1) partitioning into equal parts minimizing the number of connections between parts and 2) clustering using modularity or community detection. Hierarchical approaches are used for more than two parts/clusters. "Proof of principle" results and comparisons are shown for example benchmark graphs and small material systems on the simulator and D-Wave machine. The DM, ToQ, SAPI, and QBSOLV tools were used in this work.

Generative Modeling for Machine Learning on the D-Wave Sunil Thulasidasan

Abstract:

We will discuss training a generative machine learning model on the D-Wave. This model, known as a Restricted Boltzmann Machine, is often used as a building block for Deep Learning Systems because of its ability to learn features in an unsupervised way. Training such models involve sampling from a Boltzmann distribution at each step, which in theory is achieved by running a Markov Chain to convergence. This is the computational bottleneck in such systems and here we will explore the possibility of using the D-Wave — which is a physical Boltzmann machine — to accelerate this process



Presentation Abstracts

by using the statistical properties of the energy distribution of states in the D-Wave. We will compare the generative ability of D-Wave to classical methods for a data set of hand-written digits.

Learning Sparse Representations for Object Classification Using the

Quantum D-Wave 2X Machine.....Garrett Kenyon, Nga T. T. Nguyen

Abstract:

Solving for sparse representations is an important problem in computer vision, with applications including denoising, upsampling, compression and object detection. On conventional digital computers, it is customary to employ an L1 sparseness penalty, which results in a convex cost function that possesses a single global minimum. However, an L1 sparseness penalty artificially suppresses large activations. Ideally, we would like to apply an L0 sparseness penalty, which assigns a fixed cost to all non-zero activation coefficients, but this results in an extremely non-convex cost function that possesses multiple local minima. Indeed, sparse coding with an L0 sparseness penalty falls into an *NP-hard* complexity class of decision problems. *Here, we present solutions for sparse representations under the constraint of an L0 sparseness penalty using quantum annealing on the D-Wave 2X.* Quantum annealing provides a strategy for finding sparse representations that correspond to good local minima of a non-convex cost function. We use a dictionary of randomly-generated features vectors that are consistent with the coupling constraints imposed by the D-Wave 2X architecture, including defects resulting from missing qubits and couplers. The restriction to nearest-neighbor couplings between qubits is consistent with the local nature of interactions between neurons. In our preliminary studies, sparse binary representations and associated classifications were obtained for approximately ~40 canny-filtered, 24x24 center-cropped CIFAR-10 images, divided into non-overlapping patches of different sizes. In future work, we hope to obtain sparse representations using the same (or similar) sets of randomly-generated features implemented on LLNL's TrueNorth system, while we will obtain corresponding simulated results using a conventional sparse solver. *For all sets of sparse representations, we intend to train a linear classifier (SVM or single layer perceptron) to classify the sparse representations of the 10,000 test images after training on the sparse representations of the 50,000 training images.* We expect that classification performance will be far superior for the sparse representations generated on the D-Wave 2X compared to the sparse representations obtained on a conventional workstation using gradient descent. Because the TrueNorth system employs spiking neurons, it is possible that the resulting sparse representations will be less confounded by local minima. If so, we hope to show that spiking neural networks in the brain mimic some aspects of quantum annealing by more comprehensively exploring the energy landscape.

Phylogenetics through Quantum Annealing..... Marcus Daniels

Abstract:

Typically, there exist alternative ways to embed an Ising or QUBO problem on a D-Wave Chimera graph. However, not all systems are uniformly sensitive to noise, and it is far from clear how to predict the consequences of errors for a given embedding. I'll discuss a programmable embedding technique in which each constraint of a problem can have an embedding structurally custom fit to the importance of each constraint. Because constraint violation often is manifested through the phenomena of chain breaking, alternative routing and reinforcement for chains is explicitly represented using a logic



Presentation Abstracts

programming language. A long term objective of this work is to learn if there are embedding tactics that are especially robust, and can tolerate hardware limitations.

D-wave Quantum Computer as an Efficient Classical Sampler Andrey Lokhov

Abstract:

In this presentation we report two developments related to the generation of samples by the D-wave computer and to the reconstruction of the input Ising model. First, we verify the hypothesis that the D-wave quantum annealer can efficiently sample from the Boltzmann distribution corresponding to the Ising model at some established effective temperature. Second, we present an accelerated implementation of the efficient LANL "Screening" algorithm which allows to learn the underlying Hamiltonian from the D-wave generated samples.

Challenges and Successes of Solving Binary Quadratic Programming

Benchmarks on the DW2X QPU Coffrin Carleton

Abstract:

Binary Quadratic Programs (BQP) are a challenging class of NP-Hard discrete optimization problems with wide variety of real-world applications. With over 1000 qubits, the DW2X QPU is the first quantum computer with the potential to encode extremely challenging BQPs, such as those considered in the Quadratic Programming Library (QPLib). In this presentation we report our attempts to use the DW2X to solve cases from the QPLib and highlight some core challenges in using the DW2X hardware and software tools on a variety of real-world BQPs.

Topological Sphere Packing on the D-Wave David Nicholaeff

Abstract:

We demonstrate mapping a particular NP-hard continuous non-convex quadratic optimization problem with constraints to the discrete quadratic unconstrained binary optimization programming model of the adiabatic D-Wave quantum computer. Spheres of heterogeneous radii are packed into a discretized grid with bins enumerated via space-filling curves. The facility mirrors an algorithm for the bin packing problem and utilizes sequential unconstrained minimization techniques with topology as the parameter.

Quantum Uncertainty Quantification for Physical Models using ToQ.jl..... Dan O'Malley

Abstract:

We demonstrate ToQ.jl, a high-level programming language for the D-Wave based on Julia. We show how to use ToQ.jl with qbsolv to solve relatively large problems with the D-Wave. Next, we show some preliminary results for using the D-Wave as a sampling machine to do uncertainty quantification (UQ) for a diffusion model. This UQ approach treats samples from the D-Wave machine as coming from a Boltzmann distribution and uses importance sampling to estimate statistics for a target distribution.

